Zero+ Intelligence Models in Financial Markets

The traditional paradigm in economics is one of rational utility maximizing agents. Recognizing limitations in human cognition, economists have increasingly explored models in which agents have bounded rationality.

While no one would dispute the fact that agents in financial markets behave strategically however there are some problems where other factors may be more important. Previous work along these lines includes Becker's work, where it showed that random agent behavior and a budget constraint are sufficient to guarantee the proper slope of supply and demand curves. Gode and Sunder demonstrated that if one replaces the students in a standard classroom economics experiment by zero-intelligence agents with a budget constraint they perform surprisingly well.

Recently a model of a continuous double auction it has been developed [2]. This model based on the simple assumption that agents, subject to constraints imposed by current prices, place orders to buy or sell at random. In these two pages I am going to show how for certain problems such an approach, also known as "zero intelligence", can make surprisingly good quantitative predictions. A validation of this model has been done using real data from the London Stock Exchange.

The continuous double auction is the most widely used method of price formation in modern financial markets. The auction is called "double" because traders can submit orders both to buy and to sell and it is called "continuous" because they can do so at any time. An order that does not cross the opposite best price, and so does not result in an immediate transaction, is called a *limit order* (LO). An order that does cross the opposite best price, and thus causes an immediate transaction, is called a *market order* (MO). Buy and sell limit orders accumulate in their respective queues, while buy and sell MOs cause transactions that remove LOs. A LO can also be removed from its queue by being canceled. The lowest selling price offered at any point in time is called the *best ask*, $a(t)$, and the highest buying price the *best bid*, $b(t)$. The bid-ask *spread* $s(t) \equiv a(t) - b(t)$ measures the gap between them.

The model assumes that two types of agents place orders randomly according to independent Poisson processes, as shown in Fig. 1. Impatient agents place MOs randomly with a Poisson rate of μ shares per unit time. Patient agents, in contrast, place LOs randomly in both price and time. Buy LOs are placed uniformly anywhere in the semi-infinite interval $-\infty < p < a(t)$, sell LOs anywhere in $b(t) < p < \infty$. Both buying and selling LOs arrive

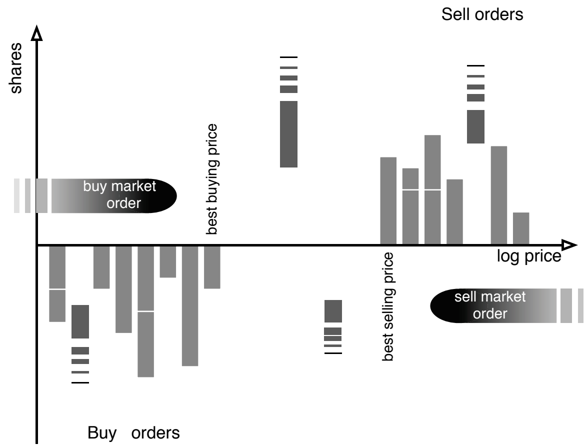

Figure 1: A random process model of the continuous double auction. Stored LOs are shown stacked along the price axis. New LOs are visualized as randomly falling down, and new buy orders as randomly "falling up". Limit orders can be removed spontaneously or they can be removed by MOs of the opposite type.

with same Poisson rate density α, which is measured in shares per unit price per unit time. The log-price p is continuous and is independent of arrival time. Both limit and market orders are of constant size σ (measured in shares). Queued LOs are canceled according to a Poisson process with a fixed rate δ per unit time. All of these processes are independent except for indirect coupling through the boundary condition: as new orders arrive they may alter the best prices $a(t)$ and $b(t)$, which in turn changes the boundary conditions for subsequent limit order placement. It is this feedback between order placement and price diffusion that makes this model interesting. The mean value of the spread predicted based on a mean field theory analysis of the model is

$$\hat{s} = (\mu/\alpha) f(\sigma\delta/\mu). \tag{1}$$

The scaling law above is reasonable in that it predicts that the spread increases when there are more MOs or cancellations (which remove stored LOs), and decreases with more LOs (which fill the spread in more quickly).

Another prediction of the model concerns the price diffusion rate, which drives the volatility of prices and is the primary determinant of financial risk. If one assumes that prices make a random walk, then the diffusion rate measures the size and frequency of its increments. Numerical experiments indicate that the short term price diffusion rate is to a very good approximation given by the formula

$$\hat{D} = k\mu^{5/2}\delta^{1/2}\sigma^{-1/2}\alpha^{-2}, \tag{2}$$

where k is a constant. This formula is reasonable in that it predicts that volatility increases with LO removal (either by MOs or by cancellations) and decreases with LO placement. The model is tested cross-sectionally over eleven stocks [1]. For each stock its average order flow rates are measured and the predicted average spread $\hat{\sigma}$ and diffusion rate $\hat{D}$ for that stock are calculated using equations 1 and 2. These predicted values are compared to the actual values of the spread $\bar{\sigma}$ and diffusion rate $\bar{D}$ measured from the data. The comparison is done via linear regressions of the predicted against the actual values.

Measurement of the parameters μ and σ is straightforward, however a problem occurs in measuring the parameters α and δ because in the real data LO placement and cancellation are concentrated near the best prices . An auxiliary assumption can be made: the order placement is uniform inside a price window around the best prices. To test this hypothesis a regression of the form $\log \bar{s} = A \log \hat{s} + B$is performed. In Fig. 2 the predictions are compared to the actual values. The least squares regression, shown also in Fig. 2, gives $A = 0.99 \pm 0.10$ and $B = 0.06 \pm 0.29$ with an $R^2 = 0.96$, proving that the model explains most of the variance.

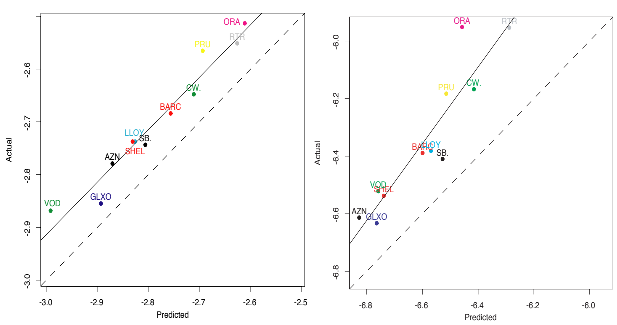

Figure 2: Regressions of predicted values based on order flow using equation 1 versus actual values for the log spread (left) and for the log diffusion rate (right).

As for the spread, the predicted price diffusion rate based on order flows is compared to the actual price diffusion rate $\bar{D}_i$ for each stock averaged over the 21 month period (Fig. 2).

The market impact is important because it provides a probe of the supply and demand functions in the order book. When a MO of size ω arrives, if it removes all LOs at the best bid or ask it will immediately change the midpoint price $m \equiv (a + b)/2$. We define the average market impact function as $\phi(\omega) = E[\Delta p|\omega]$ where Δp is the difference between the price just before and after a MO. Normalizing the price shift and the order size by appropriate dimensional scale factors based on the daily order flow rates, the collapse of the market impact can be plotted on an universal curve. This curve (Fig.3) reveals an universal supply and demand function. The market impact collapse illustrates that the non-dimensional coordinates dictated by the model provide substantial

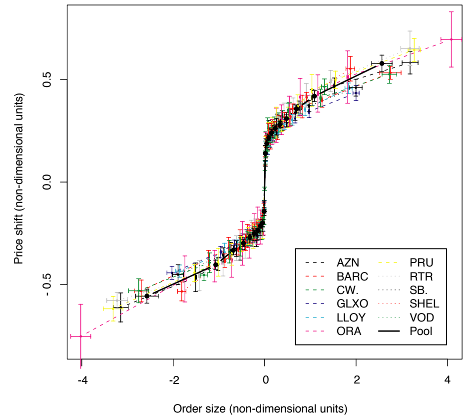

Figure 3: The average market impact as a function of the mean order size. The price differences and order sizes for each transaction are normalized by the non-dimensional coordinates.

explanatory power and reveals a universal supply and demand function. These results are remarkable because the underlying model largely drops agent rationality. At the same time the model fails to captures many important features of the real markets like market efficiency. A way to improve the model is to incrementally add new assumptions about agents behavior. It has been shown that liquidity variations are driving the large price fluctuations. The distribution of large price changes merely reflects the distribution of gaps in the limit order book.

I'm working on a model of the limit order placement process that is micro founded and generates a power low in LOs gap distribution. This project is motivated by some empirical evidences. Psychological experiments suggested that most people do not use exponential weights in considering future utility and a power law provides a better fit to the empirical data. For example in a world of uncertain interest rates, the loss of utility with time must be weighted by the distribution of interest rates: $u(t) = \int P(r)e^{-rt}dr$. If P(r) is distributed as a log-normal then $u(t)$ is a power low. Having non-zero-intelligence agents trying to maximize the expected utility under a non exponential discounting rate can be the explanations of the very important power low in the stock returns.

References

[1] J. D. Farmer, P. Patelli, and I. Zovko. The Predicitive Power of Zero Intelligence in Financial Markets. *PNAS*, 102 (6) pp. 2254-2259, 2005.

[2] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy. Statistical theory of the continuous double auction. *Quantitative Finance*, 3(6) pp. 481–514, 2003.

Contact Information: Paolo Patelli – CNLS and T-13, Los Alamos National Laboratory, MS-B213, Los Alamos, NM, 87545, Ph.: (505)667-9468, email: paolo@lanl.gov.